Rhythm Demand Planner™
Statistical Methods

WP-168

Purpose

This document outlines the statistical methods most frequently used in Rhythm Demand Planner.

Effective forecasting procedures using Rhythm Demand Planner are often built to specifically encourage and facilitate appropriate users to add their knowledge of future product, customer, company, market, and environmental expectations as well as of historical data and events before the forecasts are considered to be final.

Exponential Smoothing

Exponential smoothing is one of the most frequently used statistical forecasting methods. This technique was originally popularized by Robert G. Brown, Charles C. Holt and Peter R. Winters.[1] Its popularity is a result of 1) its ability to generate forecast accuracy that has been hard to beat with statistical methods alone, even using much more complicated models, 2) its ability to deal intelligently and flexibly with situations like forecasting with very limited historical data stored or even available, and 3) the fact that the technique is relatively easy to implement.

When exponential smoothing methods are used in Rhythm Demand Planner, they are typically implemented with rule based logic that dynamically reacts to key characteristics of each time series. An expert system style approach is then used to dynamically determine the shape of the forecast that appears to be most appropriate for each time series each time the forecasts are updated.

Single Exponential Smoothing

Calculates a moving average that gives the most recent period a weight defined by a smoothing constant that is set between zero and one. The most recent value of the moving average is given a weight equal to one minus this smoothing constant. All future forecasts have the same value.

"Double" Exponential Smoothing

Decomposes a historical time series into a level and a trend, using one smoothing constant to calculate the level as above, and a second constant to calculate the trend, which is defined as the average change in the level. Level and trend are then combined to generate a forecast that extends forward over the time horizon of interest. In its simplest form, this method generates forecasts that graph as a straight line.

"Triple" Exponential Smoothing

Decomposes a historical time series into a level, a trend, and a set of seasonal factors. We use separate smoothing constants for level, trend, and seasonality. These three components are then combined to generate forecasts for the time horizon of interest. In its simplest form, this method's forecasts move above and below a straight trend line in a pattern that repeats with, for example, an annual cycle.

"Triple" Exponential Smoothing plus Trend Curvature Extrapolation

Decomposes as above (level, trend, and seasonal factors) with the addition of a Trend Curvature factor that reflects the average rate of change of the trend. When these four averages are combined to generate a forecast, the result will move above and below a curved trend line in a pattern that, as in the previously described method, repeats with, for example, an annual cycle. This curved trend line can be thought of as an extrapolation of the way the time series is turning up or turning down.

"Triple" Exponential Smoothing plus Causal Factors — With or Without Trend Curvature

This method decomposes a historical time series into 4 components to estimate the effects of events, e.g.: promotions, price changes, and/or competitive activity), in addition to level, trend, and seasons. When the components are combined to generate a forecast, the shape of the forecast will depend on the number and timing of future causal events that are entered as data.

Exponential Smoothing in Rhythm Demand Planner

Rhythm Demand Planner implements exponential smoothing in particular (and forecasting, in general, where applicable) with formulas and/or procedures that specifically address many special conditions in the data. For example, Rhythm Demand Planner will:

- Use transformations on the time series to improve forecast accuracy. A logarithmic transformation to deal with ratio changes in the data over time is one example of this technique.
- Support separation of data into, for example, a "base" or "turn" level and a set of "incremental" or "caused" values using any method consistent with available data.
- Apply trend, seasonal, and causal calculations as either additive or multiplicative factors.

- Specifically deal with making the best possible use of limited historical data, including situations with new products or customers. "Start up values" are calculated to produce reasonable forecasts as soon as possible.
- Specifically deal with the most appropriate time period definitions for each application, including any variety of fiscal calendar that the organization might use.
- Use any seasonal factors that are available from other sources, and replace them as appropriate with factors that reflect seasonal variations observed at any level of detail.
- Smooth observed seasonality as appropriate, if, for example, the historical data is in weekly time periods with large week to week variations that mask the underlying seasonality.
- Establish zero, or any other user defined number as a minimum forecast, and any user defined number as a maximum forecast, when a minimum and/or maximum is appropriate.
- Avoid the forecast instability that can be associated with unrestricted extrapolation of either trends or trend curvatures, explosive percent calculations when small numbers appear in denominators, or the affect of occasional historical data points that seem to be out of context with the general pattern of the data.
- Specifically identify statistically definable outliers based on either user specified or default criteria, and either ignore the observation completely, or replace it with a value defined by an appropriate set of rules.
- Reflect planned transitions from existing products to new ones.
- Support bottom up, top down, and/or middle out forecasts.
- Allow authorized users to specify smoothing constants and forecast method selection for any specific time series as an override to the automatic selections built into the system.

Other Statistical Methods

Ater exponential smoothing, the most frequently used statistical methods are listed below.

Intermittent Demand Modeling

Based on the approach published by J.D. Croston in 1972, this method calculates the average number of periods between demands, and the average size of demands as a function of the number of periods since the last demand. Future demands are then forecast to take place at this calculated, most likely interval in an amount equal to the average demand associated with this interval. Special treatment is used for the first future demand forecast whenever the average interval has already passed. Variations on this theme are used for "Lumpy Demand".

Simple Moving Average

Uses the last "n" periods (specified by user) to forecast the next and future periods. All future forecasts have the same value.

Single Linear Regression

Performs a "least square" fit of the time series data. The result is a straight line that minimizes the sum of the squared values of the differences between the line and the individual values of the time series data.

Multiple Regression

Calculates a "least square" fit of user specified independent variables (causal factors) to the dependent variable (the time series being forecasted). The result is expressed as a formula that generates an estimate of the time series being forecasted based on a constant term plus the sum of the result of multiplying a (coincident) value for each independent variable by a multiplier (or coefficient) that the method derives for each causal factor. This regression technique finds the value for the constant and for the causal factor multipliers that best fits the result of the formula to the time series being forecasted. The measure of the quality of the fit is called its correlation coefficient, usually referred to as its "R squared". The shape of the forecast is completely determined by the future values of the causal factors.

Auto Regression

Uses single linear regression or multiple linear regression with prior values of the dependent variable used as independent variables. The result is an equation as described in the multiple regression method above, and forecasts take on a repetitive cyclical shape as short term forecasts are used as the basis for longer range forecasts.

Non-Statistical Methods

Non-statistical methods are fully supported by Rhythm Demand Planner in an OLAP (On Line Analytical Processing), and/or an EIS (Executive Information System) environment.

Replication of Prior Year Actuals — Plus or Minus Amount or Percent

This is just one example of a potentially unlimited list of non-statistical methods. Rhythm Demand Planner's Excel or Lotus 1-2-3 like spreadsheet and graphic user interface, plus a set of time series oriented, pre-programmed computations, and a continually expanding set of user functions available to define spread sheet rows with point and click operations, enable even casual users to define their own basis for generating or evaluating forecasts.

Forecast Accuracy

Rhythm Demand Planner typically picks a single forecast to suggest based on customized historical accuracy criteria. This functionality is referred to as "Pick Best". When Rhythm Demand Planner looks at the recent forecast accuracy of a number of alternate forecasting methods as its criteria to select the method most likely to generate the best future accuracy for any given time series, it can use any measure of forecast accuracy. The forecasting method selected, and, therefore, the forecast generated, can vary widely depending on the measure of forecast accuracy used as the method selection criteria.

The best measure of forecast accuracy for evaluating alternate forecasting methods, is usually the same as the measure used to monitor the quality of forecasting results.

Rhythm Demand Planner tracks forecast accuracy using the full range of mathematical and statistical measures we have encountered in practice or in forecasting literature. We consider it essential to include the capability to reproduce any measure used in practice so that the people and organizations converting to Rhythm Demand Planner can continue to track the performance of their forecasts on a consistent basis, and use this track record to support their continuous forecast improvement programs. An illustrative list of forecast accuracy measures frequently used in Rhythm Demand Planner follows.

Error

(Actual minus Forecast) or (Forecast minus Actual). The calculation can use one period, or the sum of any user definable number of periods. The period or periods used can be the first "n" periods of the forecast horizon, or any other user defined part of the forecast horizon.

Mean Error

A moving average of Error, as defined above. It can use a simple arithmetic average of a user definable number of periods, an exponentially smoothed moving average of Error with a user defined smoothing constant, or any other set of user defined weights. This measure is often referred to as the Bias of the forecast.

Percent Error

Error, as defined above, stated as a percent of Actual or as a percent of Forecast.

Mean Percent Error

A moving average of Percent Error, as defined above. It can use a simple arithmetic average of a user definable number of periods, an exponentially smoothed moving average of Percent Error with a user defined smoothing constant, or any other set of user defined weights.

Absolute Error

Actual minus Forecast as a positive number without regard to the sign of the result. The calculation can use one period, or the sum of any user definable number of periods. The period or periods used can be the first "n" periods of the forecast horizon, or any other user defined part of the forecast horizon.

Mean Absolute Error

A moving average of Absolute Error, as defined above. It can use a simple arithmetic average of a user definable number of periods, an exponentially smoothed moving average of Mean Absolute Error with a user defined smoothing constant, or any other set of user defined weights.

Absolute Percent Error

Percent Error, as defined above, stated as a positive number without regard to the sign of Percent Error.

Mean Absolute Percent Error

A moving average of Percent Error, as defined above. It can use a simple arithmetic average of a user definable number of periods, an exponentially smoothed moving average of Mean Absolute Percent Error with a user defined smoothing constant, or any other set of user defined weights.

Squared Error

The square of (Actual minus Forecast). The calculation can use one period, or the sum of any user definable number of periods. The period or periods used can be the first "n" periods of the forecast horizon, or any other user defined part of the forecast horizon.

Mean Squared Error

A moving average of Squared Error, as described above. It can use a simple arithmetic average of a user definable number of periods, an exponentially smoothed moving average of Error with a user defined smoothing constant, or any other set of user defined weights.

Root Mean Squared Error

The square root of Mean Squared Error as defined above.

Standard Error of Estimate

The square root of {[the sum of n observations of the square of (actual minus forecast)] divided by (n minus1)}.

The Standard Error of Estimate can be approximated by multiplying the comparable Mean Absolute Error, as defined above, by the constant 1.25. This measure is often used to estimate statistical confidence intervals around the forecast. It is also usually the best measure of forecast accuracy to be used in inventory calculations. The Standard Error of Estimate is also referred to as Standard Deviation.